

PSI-NASA
20V-51-TM
013250

**NEW SOFTWARE FOR LARGE DENSE SYMMETRIC
GENERALIZED EIGENVALUE PROBLEMS
USING SECONDARY STORAGE***

**ROGER G. GRIMES
HORST D. SIMON**

**Boeing Computer Services
Engineering Technology Applications Division
P.O. Box 24346, M/S 7L-21
Seattle, WA 98124-0346**

May, 1987

Keywords: eigenvalue problems, block Householder reduction, vector computers, super-computers, out-of-core algorithms.

AMS Subject Classification: 65F15, 88H10

1. INTRODUCTION

Quantum mechanical bandstructure computations require repeated computation of a large number of eigenvalues of a symmetric generalized eigenproblem. In previous work of the authors [5] the application of a block Lanczos algorithm, and a modification of the packed EISPACK routines, have been considered for the numerical solution of such eigenvalue problems arising in the theoretical calculation of pressure and volume at metallization of BaTe (see [10]). Both approaches were limited in the size of problem, which could be handled by the existing software. The largest feasible problem for each method on a 4 Mword Cray X-MP/24 was about of the order of 1900. Here we report numerical results with some new software, which has been developed recently. This software permits the user to fully utilize a fast secondary storage device such as the SSD. On a 128 Mword SSD now the solution of generalized eigenvalue problems of order up to 8000 is feasible. In this

*This work is supported through NSF Grant No. ASC-8519354.

note we give a brief summary of the new software and present some performance results and a comparison with the results in [5].

We consider algorithms for the efficient solution of the symmetric generalized eigenproblem

$$AX = BXA. \quad (1)$$

The new software provides a true out-of-core implementation of some standard routines from EISPACK. The approach is to first reduce (1) to standard form by applying the Cholesky factor L of B as follows

$$(L^{-1}AL^{-T})(L^T X) = (L^T X)\Lambda$$

or

$$CY = Y\Lambda, \quad (2)$$

where

$$\begin{aligned} C &= L^{-1}AL^{-T} \\ Y &= L^T X. \end{aligned} \quad (3)$$

The matrix B is factored using the block Cholesky algorithm where B is initially stored in secondary storage and is overwritten with its factor L [4,9].

The next step in the standard EISPACK approach would be to reduce C to tridiagonal form using a sequence of Householder transformations. The application of Householder transformations to reduce C to tridiagonal form requires access to all of the unreduced portion of C for the reduction of each column. This access requirement would incur $O(n^3)$ I/O transfers. This amount of I/O is too high for an out-of-core implementation, since it would effectively limit the efficiency of the algorithm by the transfer speed to and from secondary storage.

Instead a block Householder transformation is used to reduce C to a band matrix, where the band matrix will have small enough bandwidth to be held in central memory. This reduction is accomplished by first symmetrically partitioning the matrix C into equal sized blocks with maximum block size p . Then for each block column we compute the QR factorization of the matrix comprised of the blocks which are below the diagonal block. The product of Householder transformations used to compute the QR factorization is accumulated in the WY representation described in [1]. This representation is then applied simultaneously to the left and the right of the remainder of the matrix. When all block columns except the last is reduced in this fashion, the result is a band matrix with bandwidth p which is similar to the matrix C in (2). The I/O requirements of the block Householder transformation are of $O(n^3/p)$. Hence for any reasonable block size, the performance of the algorithm will no longer be I/O bound.

The banded eigenproblem is then transformed to tridiagonal form using an enhanced version of EISPACK subroutine BANDR. Finally EISPACK subroutine TQLRAT is used to compute the eigenvalues.

Specified eigenvectors are computed by first computing the associated eigenvectors of the band matrix and then back transforming them to Y using the accumulated block Householder transformations. Application of L to Y back transforms Y to X , the eigenvectors of the original problem (1). The eigenvectors of the band matrix are computed using inverse iteration implemented in a much modified version of EISPACK subroutine BANDV.

This algorithm has been implemented in a pair of subroutines HSSXGV and HSSXG1 which, respectively, compute the eigenvalues and a specified set of eigenvectors of (1). The usage of these two subroutines as well as subroutines HSSXE1 and HSSXE1, which compute the eigenvalues and specified eigenvectors of (2), is described in [7]. A more detailed description of the algorithms used is given in [6].

2. RESOURCE REQUIREMENTS

The new software makes extensive use of both the central processing unit and secondary storage of a computer. This section will discuss the requirements of the algorithm for secondary storage, amount of I/O transfer between central memory and secondary storage, storage requirements for central memory, and the number of floating point operations.

Secondary storage for this algorithm is used to hold B which is overwritten with L , to hold A which is overwritten with C and later the band matrix, and to hold the matrices W and Y . Approximately $n^2/2$ storage is used for each of the 4 matrices, thus the total amount of secondary storage required is approximately $2n^2$. A Solid-state Storage Device (SSD) with 128 million words secondary storage with fast access on a Cray X-MP would allow problems of order up to 8000 before overflowing the SSD to slower secondary storage on disks.

The amount of I/O transfer between central memory and secondary storage storing the matrices A , C and the resulting band matrix is $\frac{7}{4}n^3/p$ real words. The I/O transfer for the unit storing L , the Cholesky factor of B , is n^3/p . The total amount of I/O transfer is $\frac{11}{4}n^3/p$.

The maximum central memory requirements for the new algorithm is

$$2np + \max(4p^2, np + 3n + p, np + n + p(p + 1)),$$

where the three terms correspond to Cholesky factorization of B , reduction to banded form, and computation of the eigenvectors. A good working estimate for the amount of central memory required is $3np$. For the size of problems being considered (e.g., $n = 5000$ and $p = 50$) this is less than 1 million words of working storage.

The operation count for the computing the eigenvalues of the eigenproblem in (1) by this approach is approximately $\frac{11}{3}n^3 + 10n^2p$ + lower order terms. This operation count comes from analyzing the 4 major components of the algorithm. The operation counts for these components are as follows:

<u>Component</u>	<u>Operation Count</u>
Cholesky factorization	$\frac{1}{3}n^3$
Reduction to standard form	$2n^3$
Block Householder transformations	$\frac{4}{3}n^3 + 2n^2p$
QR factorizations	$2n^2p$
Reduction of band matrix	$6n^2p$
Total	$\frac{11}{3}n^3 + 10n^2p$

For an execution with $n = 1492$ and $p = 50$ the monitored operation count was 13.3×10^9 . The operation count computed by the above formula is 13.2×10^9 which is within 1% of the actual count.

The computational kernels used in this algorithm are Cholesky factorization, block solves with the Cholesky factor, QR factorization, matrix multiplication, and a banded eigenvalue solver. All of these kernels except for those used in solving the banded eigenproblem perform well on vector computers. In fact, they have computational rates in the 150 to 190 megaflop range on a CRAY X-MP for the problem sizes being discussed in this application. The code is portable and can be implemented efficiently on other vector supercomputers, whenever high performance implementations of the Level 2 BLAS [2] are available.

3. PERFORMANCE RESULTS

A parameter study for the optimal choice of blocksize, p , was performed. A choice of 63 for the block size appears to be optimal on the CRAY X-MP for the size of problems considered here. If a large number of eigenvectors will be computed a smaller blocksize might be in order.

Table 1 compares the combination of HSSXGV and HSSXG1 with the performance of the symmetric packed storage version of EISPACK and the block Lanczos algorithm described in [5] on the same eigenproblems in that paper. HSSXGV and HSSXG1 used block size 63 for all problems. The statement of the eigenproblem is to compute all eigenvalues and eigenvectors in the interval $[-1.0, 0.30]$.

n	Number of Eigenvectors	EISPACK	HSSXGV and HSSXG1	Block Lanczos
219	22	1.37	1.07	1.37
667	32	14.73	12.33	8.25
992	35	37.33	35.38	17.00
1496	35	108.54	88.74	47.70

Table 1.
Comparison of Performance with In-core Methods.
(Execution Times in Seconds)

HSSXGV and HSSXG1 not only allow larger problem sizes than the symmetric packed storage EISPACK path but have an 18% performance increase because of the block nature of the computations. Block Lanczos is indeed faster than HSSXGV and HSSXG1 for these eigenproblems. However, if more eigenvectors are required, HSSXGV and HSSXG1 is more efficient than block Lanczos. For the largest problem above, $n = 1496$, HSSXGV would become the algorithm of choice if 80 or more eigenvectors are required. Also HSSXGV computes all the eigenvalues and assumes no prior knowledge of the spectrum. Both Lanczos and the EISPACK path require an interval as input, and procede to compute all the eigenvalues in the interval. Hence HSSXGV in Table 2 has computed additional information.

The overall computational rate for the problems in Table 1 is given in Table 2 below. The highest performance with 161 MFLOPS was obtained on the largest problem. This number is quite remarkable since a significant amount of computation in the tridiagonal eigensolver is carried out in scalar mode. Assuming this computational rate the new software should be able to compute all the eigenvalues of an 8000 by 8000 dense generalized eigenvalue problem in about 3.3 hours on a Cray X-MP/24 with 128 Mword SSD.

n	HSSXGV MFLOPS	HSSGX1 MFLOPS
219	74	51
667	127	66
992	128	72
1496	161	80

Table 2.
Performance in MFLOPS.

4. SUMMARY

Software for the out-of-core solution of the symmetric generalized eigenproblem has been developed and implemented. Because of its block nature, this software is more efficient on vector computers than a related in-core algorithm. If the number of required eigenpairs is large, in particular in applications, where all eigenvalues and vectors are required, the new software is more efficient than a previous code of the authors [5] based on the Lanczos algorithm. Most importantly, the new software allows efficient solution of problems too large to fit in central memory thus providing an important computational tool to the researches in quantum mechanics and other disciplines which generate large symmetric generalized eigenproblems.

References.

- [1] Bischof, C. and Van Loan, C., *The WY Representation for Products of Householder Matrices*, SIAM J. Sci. Stat. Comp. 8, 1987, pp. s2-s13.
- [2] Dongarra, J. J., Du Croz, J., Hammarling, S., and Hanson, R., *Extended Set of Fortran Basic Linear Algebra Subprograms*, Argonne National Laboratory Report ANL-MS-C-TM-41 (Revision 3), 1986.
- [3] Garbow, B. S., Boyle, J. M., Dongarra, J. J., and Moler, C. B., *Matrix Eigensystem Routines - EISPACK Guide Extension*, Lecture Notes in Computer Sciences, Vol. 51, Springer-Verlag, Berlin, 1977.
- [4] Grimes, R., *Solving Systems of Large Dense Linear Equations*, ETA-TR-44, Boeing Computer Services, Seattle, Washington, February 1987, submitted to Supercomputing and Its Applications.
- [5] Grimes, R., Krakauer, H., Lewis, J., Simon, H., and Wei, S., *The Solution of Large Dense Generalized Eigenvalue Problems on the Cray X-MP/24 with SSD*, J. Comp. Physics 69, Number 2, 1987, pp. 471 - 481.
- [6] Grimes, R. , Simon, H., *Solution of Large Dense Symmetric Generalized Eigenvalue Problems Using Secondary Storage*, ETA-TR-XX, Boeing Computer Services, in preparation, 1987.
- [7] Grimes, R. , Simon, H., *Subroutines for the Out-of-core Solution of Generalized Symmetric Eigenvalue Problems*, Report ETA-TR-XX, Boeing Computer Services, in preparation, 1987.
- [8] Smith, B. T., Boyle, J. M., Dongarra, J. J., Garbow, B. S., Ikebe, Y., Klema, V. C., and

Moler, C. B., *Matrix Eigensystem Routines - EISPACK Guide*, Lecture Notes in Computer Sciences, Vol. 6, Springer-Verlag, Berlin, 1976.

- [9] *VectorPak Users Manual*, Boeing Computer Services Document No. 20460-0501-R1, 1987.
- [10] Wei, S.H. and Krakauer, H., *Phys. Review Lett.* 55, 1200, 1985.

